WEST Search History

DATE: Thursday, February 06, 2003

Set Name side by side	Query	Hit Count	Set Name result set
	DB=USPT; PLUR=YES; OP=ADJ		
L15	L13 and ("point-to-point")	9	L15
L14	L13 and serial	151	L14
L13	L12 and storage device	469	L13
L12	L1 and cache	810	L12
	DB=JPAB,EPAB; PLUR=YES; OP=ADJ		
L11	(front$1end or rear$1end) near10 cache	0	L11
L10	L6 near3 director	0	L10
L9	((cache/)!.CCLS. |(near10/)!.CCLS. |(l6/)!.CCLS.)	0	L9
L8	(front$1end or rear$1end) near10 storage	13	L8
L7	L6 near3 director	0	L7
L6	(front$1end or rear$1end)	1493	L6
	DB=USPT; PLUR=YES; OP=ADJ		
L5	710/126	904	L5
L4	(((710/126)!.CCLS.))	0	L4
L3	((710/56)!.CCLS.)	251	L3
L2	L1 and (front$1end or rear$1end)	20	L2
L1	((710/56 |710/126)!.CCLS. |(711/105 |711/112 |711/114)!.CCLS. |(or/)!.CCLS.)	1793	L1

END OF SEARCH HISTORY

WEST

Generate Collection

L2: Entry 1 of 20 File: USPT Feb 4, 2003

DOCUMENT-IDENTIFIER: US 6516370 B1
TITLE: Data storage system

Detailed Description Text (3):
More particularly, a rear-end portion of the directors, here directors 20.sub.0 -20.sub.3 and 20.sub.12 -20.sub.15, (only 20.sub.0 -20.sub.3 and 20.sub.14 and 20.sub.15 being shown) is electrically connected to the bank 14 of disk drives through I/O sections, here I/O adapter cards 22.sub.0 -22.sub.3 and 22.sub.12 -22.sub.15 (only 22.sub.0 -22.sub.3 and 22.sub.14 and 22.sub.15 being shown) and a front-end portion of the directors, here directors 20.sub.4 -20.sub.11, is electrically connected to the host computer 12 through I/O adapter cards 22.sub.4 -22.sub.11. In operation, when the host computer 12 wishes to store data, the host computer 12 issues a write request to one of the front-end directors 20.sub.4 -20.sub.11 to perform a write command. One of the front-end directors 20.sub.4 -20.sub.11 replies to the request and asks the host computer 12 for the data. After the request has passed to the requesting one of the front-end directors 20.sub.4 -20.sub.11, the director determines the size of the data and reserves space in the cache memory 18 to store the request. The front-end director then produces control signals on either a high address memory bus (TH or BH) or a low memory address bus (TL, BL) connected to such front-end director depending on the location in the cache memory 18 allocated to store the data and enable the transfer to the cache memory 18. The host computer 12 then transfers the data to the front-end director. The front-end director then advises the host computer 12 that the transfer is complete. The front-end director looks up in a Table, not shown, stored in the cache memory 18 to determine which one of the rear-end directors 20.sub.0 -20.sub.3 and 20.sub.12 -20.sub.15 is to handle this request. The Table maps the host computer 12 addresses into an address in the bank 14 of disk drives. The front-end director then puts a notification in a "mail box" (not shown and stored in the cache memory 18) for the rear-end director which is to handle the request, the amount of the data and the disk address for the data. Other rear-end directors poll the cache memory 18 when they are idle to check their "mail boxes". If the polled "mail box" indicates a transfer is to be made, the rear-end director processes the request, addresses the disk drive in the bank, reads the data from the cache memory and writes it into the addresses of a disk drive in the bank 14. When data is to be read from the disk drive to the host computer 12 the system operates in a reciprocal manner.

Detailed Description Text (4):

Each one of the rear-end portion of the directors 20.sub.0 -20.sub.3, 20.sub.12 -20.sub.15 is identical in construction, portions of an exemplary pair thereof, here rear-end directors 20.sub.0, 20.sub.1, being shown in FIG. 2. Thus, each one of the rear-end portions of the directors 20.sub.0 -20.sub.3, 20.sub.12 -20.sub.15 is shown to include a central processing section, CPU X and additional resources (Flash memories, etc,) arranged as shown. The exemplary pair of directors 20.sub.0, 20.sub.1, are shown in FIG. 2 to be coupled to the bank 14 of disk drives (FIG. 1) through I/O adapter cards 22.sub.0, 22.sub.1, respectively, via a system backplane 30, as indicated in FIG. 2. As discussed above in connection with FIG. 1, both the front-end portion and rear end portion of the directors are connected to the busses TH, TL, BH, BL through the backplane 30, as indicated in FIGS. 2 and 3.

Detailed Description Text (5):
As noted above in connection with FIG. 1, a front-end portion of the directors 20.sub.4 -20.sub.11 is coupled to the host computer 12 and rear-end portion of the directors 20.sub.0 -20.sub.3 and 20.sub.12 -20.sub.15 is coupled to the disk drive bank 14. The bank 14 of disk drives has a plurality of disk drive sections 80.sub.0 -80.sub.3 and 80.sub.12 -80.sub.15 (FIG. 1) of electrically connected disk drives.

Detailed Description Text (6):
Each one of the rear-end directors 20.sub.0 -20.sub.3 and 20.sub.12 -20.sub.15 has a pair of ports; i.e. a primary port P and a secondary port S. Each one of the sections 80.sub.0 -80.sub.3 and 80.sub.12 -80.sub.15 of disk drives is connected to the primary port P and to the secondary port S of a pair of the rear-end directors, as indicated in FIGS. 1 and 2, for an exemplary pair of the rear end portions of the directors, here the pair of rear-end directors 20.sub.0 and 20.sub.1. Thus, for redundancy, two different rear-end directors are able to communicate with any one sections of disk drives, as shown. Thus, for example, section 80.sub.0 (FIGS. 1 and 2) is coupled to the primary port P of rear-end director 20.sub.0 and, for redundancy in case there is a failure in director 20.sub.0, section 80.sub.0 is also connected to the secondary port S of director 20.sub.1. Likewise, section 80.sub.1 (FIGS. 1 and 2) is coupled to the primary port P of rear-end director 20.sub.1 and, for redundancy in case there is a failure in director 20.sub.1, section 80.sub.1 is also connected to the secondary port S of director 20.sub.0. Further, as shown in FIG. 1, directors 20.sub.0 and 20.sub.1 are on different sets of busses (i.e., busses TH and BL for director 20.sub.0 and busses TL and BH for director 20.sub.1) for additional redundancy.

Current US Cross Reference Classification (4):
711/112

Current US Cross Reference Classification (5):
711/114

WEST

Generate Collection

L2: Entry 2 of 20 File: USPT Dec 10, 2002

DOCUMENT-IDENTIFIER: US 6493795 B1
TITLE: Data storage system

Drawing Description Text (5):
FIG. 3 is a block diagram of a rear-end director according to the invention, such director being adapted for use in the system of FIG. 1;

Drawing Description Text (6):
FIG. 4 is a block diagram of a front-end director according to the invention, such director being adapted for use in the system of FIG. 1;

Drawing Description Text (7):
FIG. 5 is a block diagram of a dual-write force HIGH/LOW controller according to the invention and adapted for use in the front-end and rear-end directors of FIGS. 3 and 4;

Detailed Description Text (4):
More particularly, a rear-end portion of the directors, here directors 20.sub.0 -20.sub.3 and 20.sub.12 -20.sub.15, is electrically connected to the bank 14 of disk drives through I/O adapter cards 22.sub.0 -22.sub.3 and 22.sub.12 -22.sub.15, respectively, and a front-end portion of the directors, here directors 20.sub.4 -20.sub.11, is electrically connected to the host computer 12 through I/O adapter cards 22.sub.4 -22.sub.11, respectively. It should also be noted that each end of the busses TH, TL, BH, BL is terminated with a pair of master and slave arbiters, not shown, described in detail in connection with the above-referenced co-pending patent application.

Detailed Description Text (7):
Considering now the general operation of system, 10, when the host computer 12 wishes to store data, the host computer 12 issues a write request to one of the front-end directors 20.sub.4 -20.sub.11 to perform a write command. One of the front-end directors 20.sub.4 -20.sub.11 replies to the request and asks the host computer 12 for the data. After the request has passed to the requesting one of the front-end directors 20.sub.4 -20.sub.11, the director determines the size of the data and reserves space in the system memory 18 to store the request. The front-end director then produces control signals to arbitrate for the given bus on either a high address memory system bus (TH or BH) or a low address memory system bus (TL, BL) connected to such front-end director depending on the location in the system memory 18 allocated to store the data and enable the transfer to the system memory 18. The host computer 12

then transfers the data, as bursts of data, to the front-end director. The front-end director then advises the host computer 12 that the transfer is complete. The front-end director looks up in a Table, not shown, stored in the system memory 18 to determine which one of the rear-end directors 20.sub.0 -20.sub.3 and 20.sub.12 -20.sub.15 is to handle this request. The Table maps the host computer 12 address into an address in the bank 14 of disk drives. The front-end director then puts a notification in a "mail box" (not shown and stored in the system memory 18) for the rear-end director which is to handle the request, the amount of the data and the disk address for the data. Other rear-end directors poll the system memory 18 when they are idle to check their "mail boxes". If the polled "mail box" indicates a transfer is to be made, the rear-end director processes the request, addresses the disk drive in the bank, reads the data from the system memory 18 and writes it into the addresses of a disk drive in the bank 14. When data is to be read from the bank 14 of disk drives to the host computer 12 the system 10 operates in a reciprocal manner.

Detailed Description Text (10):
Each one of the rear-end portion of the directors 20.sub.0 -20.sub.3 and 20.sub.12 -20.sub.15 is identical in construction, an exemplary one therefor, here rear-end director 200 being shown in FIG. 3 to include a pair of central processing sections, CPU X and CPU Y, a dual port random access memory (RAM) section for storing the bursts of data and which includes a dual port RAM Y coupled to the Y CPU and a dual port RAM X coupled to the X CPU, shared resources (Flash memories, for, among other things, storage of programs the flow diagrams some of which are shown in FIGS. 6, 7 and 8, etc.), a pair of backplane interfaces 22H, 22L, a pair of data latches 24H, 24L, a pair of address latches 26H, 26L, a pair of I/O backplane interfaces 28.sub.1, 28.sub.2, and a dual-write force HIGH/LOW control section 30 (to be described in detail in connection with FIG. 5) all arranged as shown. The X and Y CPU sections are coupled to the bank 14 of disk drives (FIG. 1) through I/O adapter card 220 (FIG. 1) via an I/O backplane section interface 28.sub.1, 28.sub.2, as indicated. It should be noted that the director has a primary output port, P, and a secondary output port, S. As described in detail in connection the above-referenced co-pending patent application, the primary port P is connected to both I/O backplane interface 28.sub.1 and I/O backplane interface 28.sub.2. Likewise, the secondary port S is connected to both I/O backplane interface 28.sub.2 and I/O backplane interface 28.sub.1.

Detailed Description Text (13):
Each one of the front-end portion of the directors 20.sub.4 -20.sub.11 is identical in construction and is substantially the same as the rear-end directors described above in connection with FIG. 3. Thus, an exemplary one thereof, here director 20.sub.4 is shown in detail in FIG. 4 with equivalent elements being designated with the same numerical designation as that used in FIG. 3. Thus, the exemplary front-end director 20.sub.4 is shown to include a pair of central processing sections (i.e., CPU X and CPU Y), a dual port random access memory (RAM) section which includes a dual port RAM Y coupled to the Y CPU and a dual port RAM X coupled to the X CPU, shared resources (Flash memories, etc.), a pair of backplane interfaces 22H, 22L, a pair of data latches 24H, 24L, a pair of

address latches 26H, 26L, a pair of I/O backplane interfaces 28.sub.1, 28.sub.2, and a dual-write force HIGH/LOW control section 30 (to be described in detail in connection with FIG. 5), all arranged as shown. The X and Y CPU sections are coupled to the host computer 12 (FIG. 1) through I/O adapter card 224 (FIG. 1) via an I/O backplane section interface 28.sub.1, 28.sub.2, as indicated.

Detailed Description Text (14):
As with the rear-end directors, the dual port RAM section is coupled to: (1) a high address memory bus, here TH, through data latch 24H and backplane section interface 22H; and (2) a bottom low address memory bus, here BL, through data latch 24L and backplane section interface 22L, as shown. Data (i.e., burst of data) are stored in the data latches 24H, 24L in response to latch signals fed thereto by the dual-write force HIGH/LOW control section 30 on lines 32H, 32L, respectively. As noted above, section 30 is adapted to: (i) provide a global dual-write interrupt (GLB_DW) on the dual-write bus DW in response to a dual-write mode command fed to the system 10; (ii) receive a global dual-write interrupt (GLB DW) produced on the dual-write bus DW by another one of the directors 20.sub.1 -20.sub.15 and (iii) in response to the global dual-write interrupt (GLB_DW) received on the dual-write bus DW (either the GLB_DW it produced or the one produced by any one of the other directors), sequentially place the data burst stored in the dual port RAM on both of the system busses connected to such director at the same addressable location. Here, sequentially place the data burst stored in the dual port RAM on both of the system busses TH and BL at the same addressable location. Further, if the system 10 is placed in a dual-write mode, such section 30 is adapted to return to a non-dual-write mode when each of the plurality of directors 20.sub.0 -20.sub.15 is reset, in response to a reset condition provided to the system 10, to a non-dual write mode. Still further, as noted above, each one of the directors 20.sub.0 -20.sub.15 is adapted to detect a fault in the transfer of the data between the director and the memory sections 18H and 18L. In response to a detected fault between the director and one of the memory sections 18H, 18L, each one of the directors is adapted to force subsequent data transfers between the director and the other one of the memory sections.

Current US Original Classification (1):
711/114

WEST

Generate Collection

L2: Entry 4 of 20 File: USPT Nov 5, 2002

DOCUMENT-IDENTIFIER: US 6477618 B2
TITLE: Data storage system cluster architecture

Detailed Description Text (5):
FIG. 3 shows the general structure of an ICDA 14 apart from its internal array of disks. The hosts 10 and/or SANs are connected to front-end director modules (F) 20, and the internal disks of the ICDA 14 are connected to back-end director modules (B) 22. The director modules 20 and 22 are connected to a memory 24. The primary use of the memory 24 is disk data caching, wherein data that has recently been read from the disks or written by the hosts is temporarily stored in the memory 24. In the case of reads, the data for subsequent host read requests is obtained from the memory 24 rather than from a disk, resulting in significantly faster read completion times. Similarly, write requests are satisfied by initially writing the data into the memory 24 rather than to a disk. The host is allowed to continue in its processing stream, and the writes are completed to the disks later as a "background" process. Hardware and software mechanisms within the ICDA 14 are responsible for managing the data that flows through the memory 24 to satisfy several criteria, such as maintaining a coherent view of the data for the various hosts, utilizing storage space in the memory 24 efficiently, minimizing read and write delays, error handling, etc.

Detailed Description Text (6):
Although not shown in FIG. 3, the memory 24 is generally implemented as a collection of discrete memory modules, enabling flexible configuration of the ICDA 14 for different applications requiring different amounts of memory for desired performance. The director modules 20 and 22 are also discrete units that can be added or subtracted to achieve different system configurations while minimizing costs. Additionally, it is desirable to provide for redundant paths from each host 10 or SAN 16 to the various disks within each ICDA 14 to enhance system availability. It may be desired to connect each host 10 or SAN interconnect 16, for example, to multiple front-end directors 20, either of which can carry all of the data traffic if the other should fail. There may be redundant memory elements within the memory 24, and multiple paths between each director module 20, 22 and each set of such memory elements. Such redundancy can be achieved in any of a variety of ways.

Detailed Description Text (7):
Some of the above-mentioned drawbacks of the prior art are described more specifically with reference to FIGS. 1-3. While the

SAN and the ICDA 14 both exhibit a degree of scalability that provides for some flexibility in tailoring a storage system to a variety of user needs, certain aspects of this architecture impose undesirable constraints on flexibility and therefore constraints on the range of cost/performance tradeoffs that can be made by a user. One significant constraint is the memory 24 in each ICDA. The volumetric storage density of semiconductor memory is growing at a much lower rate than that of rotating storage devices such as magnetic disks. Accordingly, the maximum storage capacity of the memory 24 within an ICDA 14 may become inadequate as the storage capacity of the disks in the ICDA 14 grows over time. A similar constraint is the number of data access "ports" provided by the front-end directors 20, which likewise may become inadequate to meet system needs as more and more data is concentrated within a given ICDA 14. Another problem can arise when the desired overall storage capacity in a system is slightly greater than that provided by a single ICDA 14. In this case, each host 10 or SAN 16 having access to the data is burdened with a separate interface to an additional ICDA 14, increasing the costs associated with the system hardware.

Detailed Description Text (10):
FIG. 5 shows the general structure of an ICDA 14', apart from its internal array of disks, as used in the system of FIG. 4. The hosts 10 and/or SANs are connected to front-end director modules (F) 20', and the internal disks are connected to back-end director modules (B) 22'. The director modules 20' and 22' are connected to a switch network 34, which also has connections to a memory 24'. The switch network 34 also connects to the cluster interconnect 32.

Detailed Description Text (11):
In general, the switch network 34 is responsible for establishing connections among the participants in data transfers. When a front-end director 20' receives a request from a host 10, for example, it determines what the target of the request is, i.e., the memory 24', a back-end director 22', or the cluster interconnect 32, and directs the switch network 34 to establish the desired connection. The request, and the associated data in the case of a write, are transferred to the target. In the case of read transactions that can be satisfied from the memory 24', the connection is maintained until the desired read data has been returned to the requester. For disk reads, this first connection is broken after the request is transferred to the disk, and a new connection is established upon request of a back-end director 22' when the data has been received from the disk and is to be returned to the requesting front-end director 20'.

Detailed Description Text (17):
There are additional advantages as well. Using the cluster architecture of FIG. 4, systems can be configured in ways that are not even possible in the prior art systems. Each ICDA 14' can be tailored to provide specific functionality in the overall cluster 30 without the need for corresponding overhead. For example, certain ICDAs 141 may contain only disk storage that is made available to the other ICDAs 14' in the cluster. Such an ICDA 14' need not be configured with front-end directors 20' or even with memory 24'; the connectivity and caching functions associated with

those components can be performed in those ICDAs 14' having connections to hosts 10 and/or SANs 16. In fact, it may be desirable that one or more of the ICDAs 14' having host/SAN connections have no disks or back-end directors 22' at all. If the design of these ICDAs 14' permits, the space saved by removing the disks and back-end directors 221 can be used to increase the number front-end directors 20' and/or the size of the memory 24'. Given the growing disparity between the volumetric storage capacity of semiconductor memory and magnetic disks, it may be useful to dedicate one or more such ICDAs 14' at each connection point to the cluster 30, and utilize a smaller number of ICDAs 14' within the cluster 30 to provide the desired disk storage. Such an arrangement can be made possible, for example, by using a common interface between the switch 34 and the various components 20', 22' and 24' within the ICDAs 14' as described below.

Detailed Description Text (18):
As with the prior art ICDA 14 described above, it may be useful to incorporate redundancy features in the ICDA 14' to enhance system availability. These may include connecting each host to two front-end director modules 20', and connecting each disk to two back-end director modules 22'. In this way, each host and disk has two independent paths to the switch network 34. Additional redundancy can be achieved within the switch network 34. For example, the switch network 34 can be implemented as two separate switching sub-networks arranged in parallel, with each director module 20', 22' and the memory 24' having connections to both of the separate switching sub-networks. In the event of a failure in part or all of one of the sub-networks, data traffic can be routed through the other sub-network. It would also be desirable that each such switching sub-network have a separate interface to the cluster interconnect 32, to permit inter-ICDA transfers to continue.

Detailed Description Text (19):
Each director module 20', 22' and the memory 24' preferably has multiple independent connections, or "ports", to the switch network 34. This arrangement provides for greater parallelism by enabling a single director 20', 22' to be engaged in multiple simultaneous data transfers. At any time, the switch network 34 provides the necessary connectivity among all the participants in respective data transfers. For transfers involving the disk data cache, connections are formed between participating director modules 20', 22' and participating portions of the memory 24'. The switch network 34 may also have the ability to create connections directly between a front-end director 20' and a back-end director module 22' for certain types of transfers, bypassing the memory 24' altogether. So-called "streaming media" services, such as video delivery services, are good examples of such applications. There may be other applications in which some degree of data caching is used, but the caching is carried out by the hosts 10 using their own memory systems. Such host-based caching fundamentally alters the nature of the requests serviced by the ICDA 14', such that overall system performance may be better if the ICDA 14' performs no caching at all.

Detailed Description Text (21):
It is preferable that the directors 20', 22' and the individual

modules (not shown) of the memory 24' have substantially similar interfaces to the switch network 34. When this is the case, there is greater flexibility in configuring each ICDA 14' for the particular needs of each cluster-based system. The proportion of front-end directors 20' to back-end directors 22' can be varied as desired, as can the proportion of memory 24'. Different ICDAs 14' in the cluster 30 (FIG. 4) can provide specialized functionality. For example, it may be desirable to configure one ICDA 14' with a maximum number of disks and back-end directors 22', receiving all requests from other ICDAs 14' via the cluster interconnect 32. One or more other ICDAs 14' in such a cluster may provide mostly front-end connectivity (i.e., connections for hosts 10 and/or SANs 16) with or without memory 24' for caching. These ICDAs 14' would forward host requests via the cluster interconnect 32 to the ICDA(s) 14' in which the disk storage is concentrated.

Current US Cross Reference Classification (3):
711/114